

Ultra-M UCS 240M4單硬碟故障 — 熱交換程式 — CPS

目錄

[簡介](#)

[背景資訊](#)

[縮寫](#)

[MoP的工作流程](#)

[單HDD故障](#)

[計算伺服器上的單HDD故障](#)

[控制器伺服器上的單HDD故障](#)

[OSD-Compute伺服器上的單硬碟故障](#)

[OSPD伺服器上的單HDD故障](#)

簡介

本文檔介紹在託管Cisco Policy Suite(CPS)Virtual Network Function(VNF)的Ultra-M設定中更換伺服器中有故障的HDD驅動器所需的步驟。

背景資訊

Ultra-M是經過預打包和驗證的虛擬化移動資料包核心解決方案，旨在簡化VNF的部署。OpenStack是適用於Ultra-M的虛擬化基礎架構管理器(VIM)，包含以下節點型別：

- 計算
- 對象儲存磁碟 — 計算 (OSD — 計算)
- 控制器
- OpenStack平台 — 導向器(OSPD)

Ultra-M的高級體系結構和涉及的元件如下圖所示：




UltraM體系結構

附註：考慮Ultra M 5.1.x版本以定義本檔案中的程式。本檔案適用於熟悉Cisco Ultra-M平台的思科人員，並詳述在更換OSPD伺服器時在OpenStack層級執行的步驟。

縮寫

VNF	虛擬網路功能
ESC	彈性服務控制器
澳門幣	程式方法
OSD	對象儲存磁碟
硬碟	硬碟驅動器
固態硬碟	固態驅動器
VIM	虛擬基礎架構管理員
虛擬機器	虛擬機器
EM	元素管理器
UAS	Ultra自動化服務
UUID	通用唯一識別符號

MoP的工作流程



單HDD故障

1. 每台裸機伺服器將配置兩個HDD驅動器，以便在Raid 1配置中充當引導磁碟。在單HDD故障的情況下，由於存在Raid 1級冗餘，故障的HDD驅動器可以熱交換。
2. 請參閱以下步驟，以便在UCS C240 M4伺服器上更換故障元件：[更換伺服器元件](#)

- 3.在單硬碟發生故障的情況下，只有發生故障的硬碟會進行熱交換，因此更換新磁碟後不需要執行BIOS升級過程。
- 4.替換磁碟後，等待磁碟之間的資料同步。這可能需要幾個小時才能完成。
- 5.在基於OpenStack(Ultra-M)的解決方案中，UCS 240M4裸機伺服器可以承擔以下角色之一：計算、OSD — 計算、控制器和OSPD。
- 6.處理每個伺服器角色中的單HDD故障所需的步驟相同，本節介紹在對磁碟進行熱插拔之前要執行的運行狀況檢查。

計算伺服器上的單HDD故障

- 1.如果在充當Compute節點的UCS 240M4中觀察到硬碟驅動器出現故障，請在啟動故障磁碟的熱交換過程之前執行這些運行狀況檢查。
- 2.確定此伺服器上運行的VM，並驗證這些功能的狀態是否正常。

確定計算節點中託管的VM

確定託管在計算伺服器上的虛擬機器，並驗證它們是否處於活動狀態且正在運行。

計算伺服器包含虛擬機器的CPS VM/彈性服務控制器(ESC)組合：

```
[stack@director ~]$ nova list --field name,host | grep compute-8
| 507d67c2-1d00-4321-b9d1-da879af524f8 | VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-88a2d6fa82ea | pod1-compute-8.localdomain | ACTIVE |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c2_0_df4be88d-b4bf-4456-945a-3812653ee229 | pod1-compute-8.localdomain | ACTIVE |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-0 | pod1-compute-8.localdomain | ACTIVE |
```

附註：此處顯示的輸出中，第一列對應於通用唯一識別符號(UUID)，第二列是VM名稱，第三列是存在VM的主機名。

運行狀況檢查

- 1.登入到計算節點中託管的ESC並檢查狀態。

```
[admin@VNF2-esc-esc-0 esc-cli]$ escadm status
0 ESC status=0 ESC Master Healthy
```

- 2.登入到計算節點中託管的UAS並檢查狀態。

```
ubuntu@autovnf2-uas-1:~$ sudo su
root@autovnf2-uas-1:/home/ubuntu# confd_cli -u admin -C
Welcome to the ConfD CLI
admin connected from 127.0.0.1 using console on autovnf2-uas-1
autovnf2-uas-1#show uas ha
uas ha-vip 172.18.181.101
autovnf2-uas-1#
```

```

autovnf2-uas-1#
autovnf2-uas-1#show uas
uas version 1.0.1-1
uas state ha-active
uas ha-vip 172.18.181.101
INSTANCE IP    STATE    ROLE
-----
172.18.180.4  alive   CONFD-SLAVE
172.18.180.5  alive   CONFD-MASTER
172.18.180.8  alive   NA

```

```

autovnf2-uas-1#show errors
% No entries found.

```

3.如果運行狀況檢查正常，請繼續執行有故障的磁碟熱交換過程，並等待資料同步，因為它可能需要幾個小時才能完成。請參閱：[更換伺服器元件](#)

4.重複這些運行狀況檢查過程，以確認託管在計算節點上的虛擬機器的運行狀況已恢復。

控制器伺服器上的單HDD故障

1.如果在充當控制器節點的UCS 240M4中觀察到硬碟驅動器出現故障，請在啟動故障磁碟的熱交換過程之前執行這些運行狀況檢查。

2.檢查控制器上的起搏器狀態。

3.登入其中一個活動控制器並檢查起搏器狀態。所有服務必須在可用控制器上運行並在出現故障的控制器上停止。

```

[heat-admin@pod1-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-0 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Thu Jun 28 07:53:06 2018          Last change: Wed Jan 17 11:38:00 2018 by root
via cibadmin on pod1-controller-0

3 nodes and 22 resources conimaged

Online: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]

Full list of resources:

ip-10.2.2.2      (ocf::heartbeat:IPaddr2):      Started pod1-controller-0
ip-11.120.0.42  (ocf::heartbeat:IPaddr2):      Started pod1-controller-1
ip-11.119.0.42  (ocf::heartbeat:IPaddr2):      Started pod1-controller-2
ip-11.120.0.50  (ocf::heartbeat:IPaddr2):      Started pod1-controller-0
ip-11.118.0.48  (ocf::heartbeat:IPaddr2):      Started pod1-controller-1
ip-192.200.0.102 (ocf::heartbeat:IPaddr2):      Started pod1-controller-2
Clone Set: haproxy-clone [haproxy]
  Started: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Master/Slave Set: galera-master [galera]
  Masters: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Clone Set: rabbitmq-clone [rabbitmq]
  Started: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Master/Slave Set: redis-master [redis]
  Masters: [ pod1-controller-0 ]
  Slaves: [ pod1-controller-1 pod1-controller-2 ]
openstack-cinder-volume (systemd:openstack-cinder-volume): Started pod1-controller-0
my-ipmilan-for-controller-0 (stonith:fence_ipmilan):      Started pod1-controller-1

```

```
my-ipmilan-for-controller-1 (stonith:fence_ipmilan): Started pod1-controller-2
my-ipmilan-for-controller-2 (stonith:fence_ipmilan): Started pod1-controller-0
```

Daemon Status:

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

4. 檢查活動控制器中的MariaDB狀態。

```
[stack@director] nova list | grep control
| 4361358a-922f-49b5-89d4-247a50722f6d | pod1-controller-0 | ACTIVE | - | Running |
ctlplane=192.200.0.102 |
| d0f57f27-93a8-414f-b4d8-957de0d785fc | pod1-controller-1 | ACTIVE | - | Running |
ctlplane=192.200.0.110 |
```

```
[stack@director ~]$ for i in 192.200.0.102 192.200.0.110 ; do echo "*** $i ***" ; ssh heat-
admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_state_comment'\"; sudo mysql --
exec=\"SHOW STATUS LIKE 'wsrep_cluster_size'\"; done
```

```
*** 192.200.0.152 ***
```

```
Variable_name      Value
wsrep_local_state_comment Synced
```

```
Variable_name      Value
wsrep_cluster_size      2
```

```
*** 192.200.0.154 ***
```

```
Variable_name      Value
wsrep_local_state_comment Synced
```

```
Variable_name      Value
wsrep_cluster_size      2
```

驗證每個作用中控制器是否存在以下線路：

```
wsrep_local_state_comment: Synced
```

```
wsrep_cluster_size: 2
```

5. 檢查作用中控制器中的Rabbitmq狀態。

```
[heat-admin@pod1-controller-0 ~]$ sudo rabbitmqctl cluster_status
Cluster status of node 'rabbit@pod1-controller-0' ...
[{nodes,[{disc,['rabbit@pod1-controller-0','rabbit@pod1-controller-1',
'rabbit@pod1-controller-2']}]},
{running_nodes,['rabbit@pod1-controller-2',
'rabbit@pod1-controller-1',
'rabbit@pod1-controller-0']},
{cluster_name,<<"rabbit@pod1-controller-0.localdomain">>},
{partitions,[]},
{alarms,[{'rabbit@pod1-controller-2',[]},
{'rabbit@pod1-controller-1',[]},
{'rabbit@pod1-controller-0',[]}]}
```

6. 如果運行狀況檢查正常，請繼續執行有故障的磁碟熱插拔過程，並等待資料同步，因為它可能需要幾個小時才能完成。請參閱：[更換伺服器元件](#)

7. 重複這些運行狀況檢查過程，以確認控制器上的運行狀況已恢復。

OSD-Compute伺服器上的單硬碟故障

如果在充當OSD計算節點的UCS 240M4中觀察到硬碟驅動器故障，請在啟動故障磁碟的熱交換過程之前執行這些運行狀況檢查。

確定OSD計算節點中託管的VM

1. 計算伺服器包含ESC虛擬機器。

```
[stack@director ~]$ nova list --field name,host | grep osd-compute-1
| 507d67c2-1d00-4321-b9d1-da879af524f8 | VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-
88a2d6fa82ea | pod1-compute-8.localdomain | ACTIVE |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-
3812653ee229 | pod1-compute-8.localdomain | ACTIVE |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-
0 | pod1-compute-8.localdomain | ACTIVE |
| f5bd7b9c-476a-4679-83e5-303f0aae9309 | VNF2-UAS-uas-
0 | pod1-compute-8.localdomain | ACTIVE |
```

附註：此處顯示的輸出中，第一列對應於(UUID)，第二列是VM名稱，第三列是存在VM的主機名。

2. Ceph進程在OSD-Compute伺服器上處於活動狀態。

```
[root@pod1-osd-compute-1 ~]# systemctl list-units *ceph*

UNIT                                LOAD    ACTIVE SUB    DESCRIPTION
var-lib-ceph-osd-ceph\x2d11.mount  loaded active mounted /var/lib/ceph/osd/ceph-11
var-lib-ceph-osd-ceph\x2d2.mount    loaded active mounted /var/lib/ceph/osd/ceph-2
var-lib-ceph-osd-ceph\x2d5.mount    loaded active mounted /var/lib/ceph/osd/ceph-5
var-lib-ceph-osd-ceph\x2d8.mount    loaded active mounted /var/lib/ceph/osd/ceph-8
ceph-osd@11.service                 loaded active running Ceph object storage daemon
ceph-osd@2.service                 loaded active running Ceph object storage daemon
ceph-osd@5.service                 loaded active running Ceph object storage daemon
ceph-osd@8.service                 loaded active running Ceph object storage daemon
system-ceph\x2ddisk.slice           loaded active active   system-ceph\x2ddisk.slice
system-ceph\x2dosd.slice            loaded active active   system-ceph\x2dosd.slice
ceph-mon.target                    loaded active active   ceph target allowing to start/stop all
ceph-mon@.service instances at once
ceph-osd.target                    loaded active active   ceph target allowing to start/stop all
ceph-osd@.service instances at once
ceph-radosgw.target                loaded active active   ceph target allowing to start/stop all
ceph-radosgw@.service instances at once
ceph.target                        loaded active active   ceph target allowing to start/stop all
ceph*@.service instances at once
```

3. 驗證 OSD (HDD磁碟) 到日誌(SSD)的對映是否正常。

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph-disk list
/dev/sda :
  /dev/sda1 other, iso9660
  /dev/sda2 other, xfs, mounted on /
/dev/sdb :
  /dev/sdb1 ceph journal, for /dev/sdc1
  /dev/sdb3 ceph journal, for /dev/sdd1
  /dev/sdb2 ceph journal, for /dev/sde1
  /dev/sdb4 ceph journal, for /dev/sdf1
/dev/sdc :
  /dev/sdc1 ceph data, active, cluster ceph, osd.1, journal /dev/sdb1
/dev/sdd :
  /dev/sdd1 ceph data, active, cluster ceph, osd.7, journal /dev/sdb3
/dev/sde :
  /dev/sde1 ceph data, active, cluster ceph, osd.4, journal /dev/sdb2
/dev/sdf :
  /dev/sdf1 ceph data, active, cluster ceph, osd.10, journal /dev/sdb4
```

4. 驗證 Ceph 運行狀況和 OSD 樹對映是否良好。

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph -s
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
  1 mons down, quorum 0,1 pod1-controller-0,pod1-controller-1
monmap e1: 3 mons at {pod1-controller-0=11.118.0.10:6789/0,pod1-controller-1=11.118.0.11:6789/0,pod1-controller-2=11.118.0.12:6789/0}
  election epoch 28, quorum 0,1 pod1-controller-0,pod1-controller-1
osdmap e709: 12 osds: 12 up, 12 in
  flags sortbitwise,require_jewel_osds
pgmap v941813: 704 pgs, 6 pools, 490 GB data, 163 kobjects
  1470 GB used, 11922 GB / 13393 GB avail
    704 active+clean
client io 58580 B/s wr, 0 op/s rd, 7 op/s wr
```

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph osd tree
ID WEIGHT  TYPE NAME                UP/DOWN REWEIGHT PRIMARY-AFFINITY
-1 13.07996 root default
-2  4.35999  host pod1-osd-compute-0
  0  1.09000  osd.0                  up  1.00000  1.00000
  3  1.09000  osd.3                  up  1.00000  1.00000
  6  1.09000  osd.6                  up  1.00000  1.00000
  9  1.09000  osd.9                  up  1.00000  1.00000
-4  4.35999  host pod1-osd-compute-2
  2  1.09000  osd.2                  up  1.00000  1.00000
  5  1.09000  osd.5                  up  1.00000  1.00000
  8  1.09000  osd.8                  up  1.00000  1.00000
 11  1.09000  osd.11                 up  1.00000  1.00000
-5  4.35999  host pod1-osd-compute-3
  1  1.09000  osd.1                  up  1.00000  1.00000
  4  1.09000  osd.4                  up  1.00000  1.00000
  7  1.09000  osd.7                  up  1.00000  1.00000
 10  1.09000  osd.10                 up  1.00000  1.00000
```

5. 如果運行狀況檢查正常，請繼續執行有故障的磁碟熱交換過程，並等待資料同步，因為它可能需要幾個小時才能完成。請參閱[更換伺服器元件](#)

6. 重複這些運行狀況檢查過程，以確認在 OSD-Compute 節點上託管的 VM 的運行狀況已恢復。

OSPD伺服器上的單HDD故障

1.如果在充當OSPD節點的UCS 240M4中觀察到硬碟驅動器出現故障，建議在啟動故障磁碟的熱交換過程之前執行這些檢查。

2.檢查OpenStack堆疊和節點清單的狀態。

```
[stack@director ~]$ source stackrc
[stack@director ~]$ openstack stack list --nested
[stack@director ~]$ ironic node-list
[stack@director ~]$ nova list
```

3.從OSPD節點檢查所有底層雲服務是否處於已載入、活動和運行狀態。

```
[stack@director ~]$ systemctl list-units "openstack*" "neutron*" "openvswitch*"
```

UNIT	LOAD	ACTIVE	SUB	DESCRIPTION
neutron-dhcp-agent.service	loaded	active	running	OpenStack Neutron DHCP Agent
neutron-openvswitch-agent.service	loaded	active	running	OpenStack Neutron Open vSwitch Agent
neutron-ovs-cleanup.service	loaded	active	exited	OpenStack Neutron Open vSwitch Cleanup Utility
neutron-server.service	loaded	active	running	OpenStack Neutron Server
openstack-aodh-evaluator.service	loaded	active	running	OpenStack Alarm evaluator service
openstack-aodh-listener.service	loaded	active	running	OpenStack Alarm listener service
openstack-aodh-notifier.service	loaded	active	running	OpenStack Alarm notifier service
openstack-ceilometer-central.service	loaded	active	running	OpenStack ceilometer central agent
openstack-ceilometer-collector.service	loaded	active	running	OpenStack ceilometer collection service
openstack-ceilometer-notification.service	loaded	active	running	OpenStack ceilometer notification agent
openstack-glance-api.service	loaded	active	running	OpenStack Image Service (code-named Glance) API server
openstack-glance-registry.service	loaded	active	running	OpenStack Image Service (code-named Glance) Registry server
openstack-heat-api-cfn.service	loaded	active	running	Openstack Heat CFN-compatible API Service
openstack-heat-api.service	loaded	active	running	OpenStack Heat API Service
openstack-heat-engine.service	loaded	active	running	Openstack Heat Engine Service
openstack-ironic-api.service	loaded	active	running	OpenStack Ironic API service
openstack-ironic-conductor.service	loaded	active	running	OpenStack Ironic Conductor service
openstack-ironic-inspector-dnsmasq.service	loaded	active	running	PXE boot dnsmasq service for Ironic Inspector
openstack-ironic-inspector.service	loaded	active	running	Hardware introspection service for OpenStack Ironic
openstack-mistral-api.service	loaded	active	running	Mistral API Server
openstack-mistral-engine.service	loaded	active	running	Mistral Engine Server
openstack-mistral-executor.service	loaded	active	running	Mistral Executor Server
openstack-nova-api.service	loaded	active	running	OpenStack Nova API Server
openstack-nova-cert.service	loaded	active	running	OpenStack Nova Cert Server
openstack-nova-compute.service	loaded	active	running	OpenStack Nova Compute Server
openstack-nova-conductor.service	loaded	active	running	OpenStack Nova Conductor Server
openstack-nova-scheduler.service	loaded	active	running	OpenStack Nova Scheduler Server
openstack-swift-account-reaper.service	loaded	active	running	OpenStack Object Storage

```
(swift) - Account Reaper
openstack-swift-account.service          loaded active running OpenStack Object Storage
(swift) - Account Server
openstack-swift-container-updater.service loaded active running OpenStack Object Storage
(swift) - Container Updater
openstack-swift-container.service       loaded active running OpenStack Object Storage
(swift) - Container Server
openstack-swift-object-updater.service  loaded active running OpenStack Object Storage
(swift) - Object Updater
openstack-swift-object.service          loaded active running OpenStack Object Storage
(swift) - Object Server
openstack-swift-proxy.service           loaded active running OpenStack Object Storage
(swift) - Proxy Server
openstack-zaqar.service                  loaded active running OpenStack Message Queuing
Service (code-named Zaqar) Server
openstack-zaqar@1.service                loaded active running OpenStack Message Queuing
Service (code-named Zaqar) Server Instance 1
openvswitch.service                      loaded active exited  Open vSwitch
```

LOAD = Reflects whether the unit definition was properly loaded.

ACTIVE = The high-level unit activation state, i.e. generalization of SUB.

SUB = The low-level unit activation state, values depend on unit type.

37 loaded units listed. Pass --all to see loaded but inactive units, too.

To show all installed unit files use 'systemctl list-unit-files'.

4.如果運行狀況檢查正常，請繼續執行有故障的磁碟熱交換過程，並等待資料同步，因為它可能需要幾個小時才能完成。請參閱[更換伺服器元件](#)

5.重複這些運行狀況檢查過程，以確認OSPD節點的運行狀況已恢復。